

Random Graphons and a Weak Positivstellensatz for Graphs

LÁSZLÓ LOVÁSZ* and BALÁZS SZEGEDY

Institute of Mathematics, Eötvös Loránd University
Budapest, Hungary, and
Department of Mathematics, University of Toronto
Toronto, Ontario, Canada

February 2009

Abstract

In an earlier paper the authors proved that limits of convergent graph sequences can be described by various structures, including certain 2-variable real functions called graphons, random graph models satisfying certain consistency conditions, and normalized, multiplicative and reflection positive graph parameters. In this paper we show that each of these structures has a related, relaxed version, which are also equivalent. Using this, we describe a further structure equivalent to graph limits, namely probability measures on countable graphs that are ergodic with respect to the group of permutations of the nodes.

As an application, we prove an analogue of the Positivstellensatz for graphs: We show that every linear inequality between subgraph densities that holds asymptotically for all graphs has a formal proof in the following sense: it can be approximated arbitrarily well by another valid inequality that is a “sum of squares” in the algebra of partially labeled graphs.

Contents

1	Introduction	2
2	Preliminaries	2
2.1	Homomorphism densities and limits	2
2.2	Partially labeled graphs and quantum graphs	4
2.3	Graph parameters	4
2.4	Random graph models	5
3	Equivalent forms of the limit object	5
3.1	Graph limits and random graph limits	5
3.2	More equivalences	8
4	Weak Positivstellensatz for graphs	9

*Research sponsored by OTKA Grant No. 67867.

1 Introduction

In an earlier paper the authors proved that limits of convergent graph sequences can be described by various structures, including 2-variable symmetric, measurable functions $[0, 1]^2 \rightarrow [0, 1]$, random graph models satisfying a “consistency” and a “locality” condition, and normalized, multiplicative and reflection positive graph parameters (see Theorem 3.1 and Proposition 3.3).

In this paper we show that each of these structures has a related, relaxed version: We can drop the multiplicativity condition on the graph parameter, replacing it with the simple condition that deleting isolated nodes does not change the value of the parameter. We can drop the “locality” condition on the random graph model. We can replace the graphon by a probability distribution of the graphon. As the first main result of this paper, we prove that these relaxed versions are also equivalent.

This result will be used in adding a further equivalent structure to the list of structures describing graph limits: a probability measure on countable graphs that is ergodic with respect to the group of permutations of the nodes.

As an application, we prove an analogue of the Positivstellensatz for graphs. Many fundamental theorems in extremal graph theory can be expressed as linear inequalities between subgraph densities. For example, the Mantel–Turán Theorem is implied by the linear inequality that the density of triangles is always at least the edge-density minus $\frac{1}{2}$. (To be more precise, using “homomorphism densities” to be defined in Section 2, we get inequalities that hold true for all graphs; in terms of subgraph densities, we get in general only asymptotic results with some error terms.)

It has been observed long ago that most of these extremal results seem to follow by one of more tricky applications of the Cauchy–Schwarz inequality. We confirm this in the following sense: we show that every linear inequality between homomorphism densities that holds for all graphs can be derived, up to an arbitrarily small error term, by the Cauchy–Schwarz Inequality. To make the last phrase precise, we use graph algebras introduced by Freedman, Lovász and Schrijver in [6]. The square of an algebra element, when expanded, yields a valid linear inequality between homomorphism densities. Sums of such inequalities yield further valid linear inequalities, and our result says that such sums of squares are dense among all valid linear inequalities.

2 Preliminaries

2.1 Homomorphism densities and limits

In this paper, all graphs are simple. If we don’t quantify, we also mean that the graph is finite.

For two graphs F and G , we write $F \cong G$ if they are isomorphic, and $F \simeq G$ if they become isomorphic after their isolated nodes are deleted. So the graph U_n consisting of n isolated nodes satisfies $U_n \simeq K_0 \cong U_0$.

For two graphs F and G , let $\text{hom}(F, G)$ denote the number of homomorphisms (adjacency-preserving maps) from F to G , and $\text{inj}(F, G)$, the number of injective homomorphisms from F

to G . We consider the *homomorphism densities*

$$t(F, G) = \frac{\text{hom}(F, G)}{|V(G)|^{|V(F)|}},$$

and *subgraph densities*

$$t_{\text{inj}}(F, G) = \frac{\text{inj}(F, G)}{|V(G)| \cdot (|V(G)| - 1) \cdots (|V(G)| - |V(F)| + 1)},$$

Let \mathcal{W}_0 denote the set of symmetric measurable functions $W : [0, 1]^2 \rightarrow [0, 1]$. A *graphon* is any function in \mathcal{W}_0 . For every graph F and graphon W , we define the density of F in W by

$$t(F, W) = \int_{[0, 1]^V} \prod_{ij \in E} W(x_i, x_j) \prod_{i \in V} dx_i$$

To every graph G we can assign a graphon W_G as follows: Let $V(G) = [n]$. Split $[0, 1]$ into n intervals J_1, \dots, J_n of length $\lambda(J_i) = \alpha_i / \alpha_G$. For $x \in J_i$ and $y \in J_j$, let $W_G(x, y) = \mathbf{1}_{ij \in E(G)}$. With this construction, we have $t(F, G) = t(F, W_G)$ for all finite graphs F .

We consider on \mathcal{W}_0 the *cut norm*

$$\|W\|_{\square} = \sup_{S, T \subseteq [0, 1]} \left| \int_{S \times T} W(x, y) dx dy \right|$$

where the supremum is taken over all measurable subsets S and T , and the *cut distance*

$$\delta_{\square}(U, W) = \inf_{\phi, \psi} \|U^{\phi} - W^{\psi}\|_{\square},$$

where ϕ, ψ range over all measure preserving maps from $[0, 1] \rightarrow [0, 1]$, and $W^{\phi}(x, y) = W(\phi(x), \phi(y))$ [2, 3]. This also defines a distance between graphs by

$$\delta_{\square}(F, G) = \delta_{\square}(W_F, W_G).$$

(See [4] for more combinatorial definitions of this graph distance.)

We note that $\delta_{\square}(U, W) = 0$ can hold for two different graphons: $\delta_{\square}(W^{\phi}, W^{\psi}) = 0$ for every graphon W and measure preserving maps $\phi, \psi : [0, 1] \rightarrow [0, 1]$. (It was proved in [1] that this gives all pairs of graphons with distance 0.) We call two graphons *weakly isomorphic* if their distance is 0.

It was proved in [8] that $(\mathcal{W}_0, \delta_{\square})$ is a compact metric space.

A sequence of graphs (G_n) with $|V(G_n)| \rightarrow \infty$ is *convergent* if the densities $t(F, G_n)$ converge for all finite graphs F . This is clearly equivalent to saying that the subgraph densities $t_{\text{inj}}(F, G_n)$ converge for all finite graphs F .

It was proved in [4] that a graph sequence is convergent if and only if it is Cauchy in the δ_{\square} distance. It was proved in [9] that for every convergent graph sequence there is a limit object in the form of a function $W \in \mathcal{W}$, so that

$$t(F, G_n) \rightarrow t(F, W) \quad \text{for all graphs } F.$$

In [4] it was shown that this is equivalent to $\delta_{\square}(W_{G_n}, W) \rightarrow 0$. In [1] it was proved that this limit is uniquely determined up to weak isomorphism.

2.2 Partially labeled graphs and quantum graphs

A k -labeled graph is a graph in which k of the nodes are labeled by $1, \dots, k$ (there may be any number of unlabeled nodes). A 0-labeled graph is just an unlabeled graph. Let \mathcal{F}_k denote the set of k -labeled graphs (up to label-preserving isomorphism).

A k -labeled graph F is called *flat* if $V(F) = [k]$. Let \mathcal{F}'_k denote the set of all flat k -labeled graphs.

Let F_1 and F_2 be two k -labeled graphs. We define the k -labeled graph $F_1 F_2$ by taking their disjoint union, and then identifying nodes with the same label (if multiple edges arise, we only keep one copy). Clearly this multiplication is associative and commutative. For two 0-labeled graphs, $F_1 F_2$ is their disjoint union.

Sometimes it is more convenient to combine k -labeled graphs into a single structure. A *partially labeled graph* is a finite graph in which some of the nodes are labeled by distinct positive integers. For two partially labeled graphs F_1 and F_2 , let $F_1 F_2$ denote the partially labeled graph obtained by taking their disjoint union, and identifying nodes with the same label. Let \mathcal{F}^* denote the set of partially labeled graphs (up to isomorphism).

A *quantum graph* is defined as a formal linear combination of graphs with real coefficients. A k -labeled quantum graph is defined similarly as a formal linear combination of k -labeled graphs. The product of k -labeled graphs defined above extends to quantum graphs by distributivity: if $f = \sum_{i=1}^n \lambda_i F_i$ and $g = \sum_{j=1}^m \mu_j G_j$, then $fg = \sum_{i=1}^n \sum_{j=1}^m \lambda_i \mu_j F_i G_j$.

2.3 Graph parameters

A *graph parameter* is a real valued function defined on isomorphism types of graphs (including the graph K_0 with no nodes and edges). Let f be any graph parameter and fix an integer $k \geq 0$. We define the k -th *connection matrix* of the graph parameter f as the (infinite) symmetric matrix $M(f, k)$, whose rows and columns are indexed by (isomorphism types of) k -labeled graphs, and the entry in the intersection of the row corresponding to F_1 and the column corresponding to F_2 is $f(F_1 F_2)$. The *flat connection matrix* $M_{\text{flat}}(f, k)$ is the submatrix of $M(f, k)$ formed by rows and columns corresponding to flat k -labeled graphs (this matrix is finite).

We denote by \mathcal{M} the space of $\mathcal{F}^* \times \mathcal{F}^*$ matrices (these are infinite matrices). For a graph parameter f , we define the *full connection matrix* as the symmetric matrix $M(f) \in \mathcal{M}$, whose entry in the intersection of the row corresponding to F_1 and the column corresponding to F_2 is $f(F_1 F_2)$. Clearly this matrix contains as a submatrix all connection matrices $M(f, k)$. In the other direction, we note that every finite submatrix of $M(f)$ is contained as a submatrix in one of the matrices $M(f, k)$.

Let f be a graph parameter. We say that f is *isolate-indifferent* if $f(G) = f(G')$ whenever $G \simeq G'$. The parameter is *multiplicative* if $f(FG) = f(F)f(G)$, where FG denotes the disjoint union of the graphs F and G .

For every graph parameter f , we define its *Möbius transform* f^\dagger by

$$f^\dagger(F) = \sum_{\substack{F': V(F')=V(F) \\ E(F') \supseteq E(F)}} (-1)^{|E(F') \setminus E(F)|} f(F').$$

We say that f is *normalized* if $f(K_0) = f(K_1) = 1$. Note that for a multiplicative parameter, it would be enough to assume $f(K_1) = 1$, while for an isolate-indifferent parameter, it would be enough to assume $f(K_0) = 1$. Trivially, if a graph parameter is multiplicative and normalized, then it is isolate-indifferent.

We call a graph parameter *reflection positive* if all of its connection matrices are positive semidefinite (this is equivalent to saying that its full connection matrix $M(f)$ is positive semidefinite). We call it *flatly reflection positive* if all its flat connection matrices are positive semidefinite.

We denote by \mathcal{K} the linear space of matrices $A \in \mathcal{M}$ in which $A_{F_1, G_1} = A_{F_2, G_2}$ if $F_1 G_1 \cong F_2 G_2$, and by \mathcal{L} , the linear space of matrices $A \in \mathcal{M}$ in which $A_{F_1, G_1} = A_{F_2, G_2}$ if $F_1 G_1 \simeq F_2 G_2$. Clearly connection matrices define a bijection between matrices in \mathcal{K} and graph parameters. Under this bijection, matrices in \mathcal{L} correspond to isolate-indifferent graph parameters.

Let $\mathcal{P} \subseteq \mathcal{M}$ denote the cone of positive semidefinite matrices in \mathcal{M} . Reflection positive graph parameters correspond to matrices in $\mathcal{P} \cap \mathcal{K}$.

2.4 Random graph models

A *random graph model* is a sequence $(P_n : n = 0, 1, 2, \dots)$, where P_n is a probability distribution on graphs on $[n]$. Let \mathbf{G}_n be a random graph from distribution P_n . We say that the random graph model is *consistent*, if the distribution P_n is invariant under relabeling nodes, and if we delete node n from \mathbf{G}_n , the distribution of the resulting graph is the same as the distribution of \mathbf{G}_{n-1} .

We say that the random graph model is *local*, if for every $S \subseteq [n]$, the subgraphs of \mathbf{G}_n induced by S and $[n] \setminus S$ are independent (as random variables).

Let $\binom{\mathbb{N}}{2}$ denote the set of all unordered pairs from \mathbb{N} . Every subset of $\binom{\mathbb{N}}{2}$ can be thought of as a graph on node set \mathbb{N} , and $\{0, 1\}^{\binom{\mathbb{N}}{2}}$ is the set of all graphs on \mathbb{N} . Let \mathcal{A} denote the σ -algebra on $\{0, 1\}^{\binom{\mathbb{N}}{2}}$ generated by the sets obtained by fixing whether a given pair is connected or not.

A *random countable graph model* is a probability distribution P on $(\{0, 1\}^{\binom{\mathbb{N}}{2}}, \mathcal{A})$. Such a distribution is *consistent* if the distribution of the labeled subgraph induced by an ordered finite set S depends only on the size of S . The distribution is *local* if for any two finite disjoint subsets $S_1, S_2 \subseteq \mathbb{N}$, the subgraphs induced by S_1 and S_2 are independent (as random variables). The distribution is *invariant* if it is invariant under permutations of \mathbb{N} . The distribution is *ergodic* if there is no set $S \in \mathcal{A}$ with $0 < \pi(S) < 1$ invariant under permutations of \mathbb{N} . Invariant measures form a convex set in the linear space of all signed measures, and ergodic measures are the extreme points of this convex set.

A probability distribution on the Borel sets of $(\mathcal{W}_0, \delta_\square)$ will be called a *random graphon model*. Note that the σ -algebra of Borel sets does not distinguish weakly isomorphic graphons.

3 Equivalent forms of the limit object

3.1 Graph limits and random graph limits

We quote the following theorem, which was proved essentially in [9].

Theorem 3.1 *The following are equivalent (cryptomorphic):*

- (a) *A multiplicative, normalized graph parameter with nonnegative Möbius transform;*
- (b) *A consistent and local random graph model;*
- (c) *A consistent and local random countable graph model;*
- (d) *A graphon, up to weak isomorphism.*
- (e) *A point in the completion of the set of finite graphs with the cut-metric;*

The following theorem shows that in each of these objects, we can naturally relax the conditions, to get another important set of cryptomorphic structures.

Theorem 3.2 *The following are equivalent (cryptomorphic):*

- (a) *An isolate-indifferent, normalized graph parameter with nonnegative Möbius transform;*
- (b) *A consistent random graph model;*
- (c) *A consistent random countable graph model;*
- (d) *A random graphon model.*

Proof. We describe a cycle of constructions, mapping one object in the theorem to the next.

(a)→(b). Let f be an isolate-indifferent, reflection positive, normalized graph parameter with nonnegative Möbius transform. Using that f is isolate-indifferent, we get

$$\sum_{F: V(F)=[n]} f^\dagger(F) = f(U_n) = 1.$$

So we can construct a random graph \mathbf{G}_n on $[n]$ by

$$\mathbf{P}(\mathbf{G}_n = F) = f^\dagger(F) \quad (V(F) = [n]). \quad (1)$$

It is clear that this distribution does not depend on the labeling of the nodes. Let F_0 be a graph on $[n-1]$, and let F_0^+ be obtained from F_0 by adding n as an isolated node. Then

$$\begin{aligned} \mathbf{P}(\mathbf{G}_n \setminus \{n\} = F_0) &= \sum_{F: F \setminus \{n\} = F_0} \mathbf{P}(\mathbf{G}_n = F) = \sum_{F: F \setminus \{n\} = F_0} f^\dagger(F) \\ &= \sum_{F: F \setminus \{n\} = F_0} \sum_{F' \supseteq F} (-1)^{|E(F')| - |E(F)|} f(F') \\ &= \sum_{F' \supseteq F_0^+} f(F') \sum_{\substack{F \subseteq F' \\ F \setminus \{n\} = F_0}} (-1)^{|E(F')| - |E(F)|} \end{aligned}$$

Here the last sum is 0 unless F' contains no edges incident with the node n , and so $f(F') = f(F'')$, where $F'' = F' \setminus \{n\}$. Thus

$$\mathbf{P}(\mathbf{G}_n \setminus \{n\} = F_0) = \sum_{F'' \supseteq F_0} f(F'') (-1)^{|E(F'')| - |E(F_0)|} = f^\dagger(F_0).$$

Thus this model is consistent. We note that f can be recovered by

$$f(F) = \mathbf{P}(F \subseteq \mathbf{G}_n) \quad (V(F) = [n]). \quad (2)$$

(b)→(c). Let \mathbf{G}_n be a random graph from a consistent finite random graph model, we construct a countable random graph model by $\pi(A_F) = \mathbf{P}(\mathbf{G}_n = F) \ (V(F) = [n])$. This extends to a probability measure on the σ -algebra \mathcal{A} . It is straightforward to check that this measure is consistent.

(c)→(d). Let \mathbf{G} be a random countable graph from a consistent countable random graph model, we construct a probability distribution on the Borel sets of $(\mathcal{W}_0, \delta_\square)$. Let \mathbf{G}_n be the finite graph spanned by the first n nodes of \mathbf{G} .

We claim that with probability 1, the graph sequence (\mathbf{G}_n) is convergent. Theorem 2.11 in [4] implies that

$$\delta_\square(\mathbf{G}_n, \mathbf{G}_m) \leq \frac{10}{\sqrt{\log n}}$$

with probability $1 - \exp(-n^2/(2 \log n))$.

Let $\mathbf{H}_k = \mathbf{G}_{2^k}$, then

$$\mathbf{P}\left(\delta_\square(\mathbf{H}_k, \mathbf{H}_{k+1}) > \frac{10}{2^{k/2}}\right) < \exp\left(\frac{-2^{2^k}}{2^{k+1}}\right),$$

and so by the Borel-Cantelli Lemma,

$$\delta_\square(\mathbf{H}_k, \mathbf{H}_{k+1}) \leq \frac{10}{2^{k/2}}$$

holds for all but a finite number of values of k , with probability 1. Hence with probability 1, the sequence $(W_{\mathbf{H}_k})$ is a Cauchy sequence in $(\mathcal{W}_0, \delta_\square)$.

Now for a general value of n , let $k_n = \lceil \log \log n \rceil$. Then as before, we get that

$$\mathbf{P}\left(\delta_\square(\mathbf{G}_n, \mathbf{H}_{k_n}) > \frac{10}{\sqrt{\log n}}\right) < \exp\left(\frac{-n^2}{2 \log n}\right).$$

Again by the Borel-Cantelli Lemma,

$$\delta_\square(\mathbf{G}_n, \mathbf{H}_{k_n}) \leq \frac{10}{\sqrt{\log n}}$$

holds for all but a finite number of n , with probability 1. This proves that the sequence (\mathbf{G}_n) is Cauchy. Thus it tends to a limit graphon \mathbf{W} .

So we have described a method to generate a random graphon \mathbf{W} . For every graph F , this satisfies

$$t(F, \mathbf{W}) = \lim_{n \rightarrow \infty} t(F, \mathbf{G}_n) = \lim_{n \rightarrow \infty} t_{\text{inj}}(F, \mathbf{G}_n).$$

By the consistency of \mathbf{G} , the expectation of $t_{\text{inj}}(F, \mathbf{G}_n)$ is independent of n for $n \geq k = |V(F)|$, and so

$$\mathbf{E}(t(F, \mathbf{W})) = \lim_{n \rightarrow \infty} \mathbf{E}(t_{\text{inj}}(F, \mathbf{G}_n)) = \mathbf{E}(t_{\text{inj}}(F, \mathbf{G}_k)) = \mathbf{P}(F \subseteq \mathbf{G}_k).$$

(d)→(a). Let \mathbf{W} be a random graphon from any probability distribution on the Borel sets of $(\mathcal{W}_0, \delta_\square)$. This defines a graph parameter f by

$$f(F) = \mathbf{E}(t(F, \mathbf{W})).$$

For every fix $W \in \mathcal{W}_0$, the graph parameter $f(\cdot) = t(\cdot, W)$ is normalized, isolate-indifferent (since it is multiplicative), and has nonnegative Möbius transform (by Theorem 3.1). Trivially, these properties are inherited by the expectation. \square

3.2 More equivalences

In theorems 3.1 and 3.2, we listed several seemingly quite different objects that have turned out equivalent. In this section we show that these objects have alternative characterizations. The following characterization of graph parameters occurring in Theorem 3.1 was proved in [9].

Proposition 3.3 *Let f be a multiplicative, normalized graph parameter. Then the following are equivalent:*

- (a) f is reflection positive;
- (b) f is flatly reflection positive;
- (c) f has nonnegative Möbius transform;
- (d) $f = t(\cdot, W)$, where W is a graphon.
- (e) f is the limit of homomorphism density functions.

For graph parameters in Theorem 3.2, we have the following.

Proposition 3.4 *Let f be an isolate-indifferent, normalized graph parameter. Then the following are equivalent:*

- (a) f is reflection positive;
- (b) f is flatly reflection positive;
- (c) f has nonnegative Möbius transform;
- (d) $f = \mathbf{E}(t(\cdot, \mathbf{W}))$, where \mathbf{W} is a random graphon.
- (e) f is in the convex hull of limits of homomorphism density functions.

While the proof here is similar, there are some differences, and we include it for completeness.

Proof. (a) \Rightarrow (b) is trivial.

(b) \Rightarrow (c): The Lindström–Wilf Formula gives the following diagonalization of $M_{\text{flat}}(f, k)$: Let Z denote the $\mathcal{F}'_k \times \mathcal{F}'_k$ matrix defined by $Z_{F_1, F_2} = \mathbf{1}_{F_1 \subseteq F_2}$. Let D be the diagonal matrix with $D_{F, F} = f^\dagger(F)$. Then $M_{\text{flat}}(f, k) = Z^\top D Z$. This implies that $M_{\text{flat}}(f, k)$ is positive semidefinite if and only if $f^\dagger \geq 0$ for all graphs with k nodes.

(c) \Rightarrow (d): Let f be an isolate-indifferent, normalized graph parameter with nonnegative Möbius transform. By Theorem 3.2, it defines a random graphon \mathbf{W} such that $f = \mathbf{E}(t(\cdot, \mathbf{W}))$.

(d) \Rightarrow (e): By Theorem 3.2, each $t(\cdot, \mathbf{W})$ is the limit of homomorphism density functions for every \mathbf{W} .

(e) \Rightarrow (a): Every homomorphism density function f is reflection positive, and this is clearly inherited to their limits, and then to the convex hull of these limits. \square

The following propositions describe connections between graph-theoretic and group-theoretic properties of countable random graph models. They also indicate a connection with ergodic theory.

Proposition 3.5 *A countable random graph model is consistent if and only if it is invariant.*

Proof. It is trivial that invariant countable random graph models are consistent. Conversely, if a countable random graph model is consistent, then it defines a consistent finite graph model, which in turn defines a unique countable random graph model, independently of the labeling of the nodes. \square

Proposition 3.6 *A consistent countable random graph model is local if and only if it is ergodic.*

Proof. Let μ be an invariant probability measure on the Borel sets in $\{0, 1\}^{\binom{\mathbb{N}}{2}}$. By Proposition 3.5 it is consistent, and so by Theorem 3.2 it is defined by a random graphon. If μ is ergodic, then μ is an extreme point of all invariant distributions, and therefore this random graphon must be concentrated on a single graphon. Thus Theorem 3.1 implies that μ is local.

Conversely, if μ is not ergodic, then $\mu = \frac{1}{2}(\mu_1 + \mu_2)$, where μ_1, μ_2 are invariant probability measures and $\mu_1 \neq \mu_2$. Let \mathbf{G}_1 and \mathbf{G}_2 be random countable graphs from the distributions μ_1 and μ_2 , respectively, and let \mathbf{G} be \mathbf{G}_1 with probability 1/2 and \mathbf{G}_2 with probability 1/2. Let $S \subseteq \mathbb{N}$ be a finite set and F a labeled graph on $|S|$ nodes such that $\mathbf{P}(\mathbf{G}_1[S] = F) \neq \mathbf{P}(\mathbf{G}_2[S] = F)$. Let $T \subseteq \mathbb{N}$ be another set with $|T| = |S|$ and $T \cap S = \emptyset$. Set $a_1 = \mathbf{P}(\mathbf{G}_1[S] = F) = \mathbf{P}(\mathbf{G}_1[T] = F)$ (by invariance, these two probabilities are equal), and define a_2 analogously.

Thus we have

$$\begin{aligned} & \mathbf{P}(\mathbf{G}[S] = F, \mathbf{G}[T] = F) - \mathbf{P}(\mathbf{G}[S] = F)\mathbf{P}(\mathbf{G}[T] = F) \\ &= \frac{1}{2}(\mathbf{P}(\mathbf{G}_1[S] = F, \mathbf{G}_1[T] = F) + \mathbf{P}(\mathbf{G}_2[S] = F, \mathbf{G}_2[T] = F)) \\ &\quad - \frac{1}{4}(\mathbf{P}(\mathbf{G}_1[S] = F) + \mathbf{P}(\mathbf{G}_2[S] = F))(\mathbf{P}(\mathbf{G}_1[T] = F) + \mathbf{P}(\mathbf{G}_2[T] = F)) \\ &= \frac{1}{2}(a_1^2 + a_2^2) - \frac{1}{4}(a_1 + a_2)^2 = \frac{1}{4}(a_1 - a_2)^2 > 0. \end{aligned}$$

This shows that μ is not local. \square

4 Weak Positivstellensatz for graphs

Let $x = \alpha_1 F_1 + \dots + \alpha_r F_r$ be any quantum graph. We say that $x \geq 0$ if $t(x, W) = \sum_i \alpha_i t(F_i, W) \geq 0$ for every $W \in \mathcal{W}_0$. Hence $x \geq 0$ if and only if $\sum_i \alpha_i f(F_i) \geq 0$ for every multiplicative, reflection positive graph parameter f . Proposition 3.4 implies that this is equivalent to saying that $\sum_i \alpha_i f(F_i) \geq 0$ for every isolate-indifferent, reflection positive parameter f .

An easy example of quantum graphs $x \geq 0$ is any quantum graph of the form $\sum_i y_i^2$, where the y_i are k -labeled quantum graphs for some $k \geq 0$ (and the labels are ignored after squaring).

One may ask whether every quantum graph $x \geq 0$ can be represented this way. We don't know the answer, although based on the analogy of polynomials, the answer is probably negative. However, we prove the following weaker version, which is analogous to Lasserre's result [7] asserting that positive polynomials are approximately sums of squares.

Theorem 4.1 *Let x be a quantum graph. Then $x \geq 0$ if and only if for every $\varepsilon > 0$ there is a $k \geq 1$ and $y_1, \dots, y_m \in \mathcal{G}_k$ such that $\|x - y_1^2 - \dots - y_m^2\|_1 < \varepsilon$.*

Proof. For $n \geq k \geq 0$, let \mathcal{F}_k denote the set of k -labeled simple graphs on $[k]$ (up to isomorphism). Let Φ_k denote the operator mapping a matrix \mathcal{M} to its restriction to $\mathcal{F}_k \times \mathcal{F}_k$. Then $\mathcal{M}_k = \Phi_k \mathcal{M}$ is the space of all symmetric $\mathcal{F}_k \times \mathcal{F}_k$ matrices, and $\mathcal{P}_k = \Phi_k \mathcal{P}$ is the positive semidefinite cone in $\Phi_k \mathcal{M}$. It is also clear that $\mathcal{L}_k = \Phi_k \mathcal{L}$ consists of those matrices $A \in \mathcal{M}_k$ in which $A_{F_1, G_1} = A_{F_2, G_2}$ whenever $F_1 G_1 \simeq F_2 G_2$. We set $\mathcal{R}_k = \Phi_k \mathcal{P} \cap \Phi_k \mathcal{L}$. Clearly,

$$\Phi_k(\mathcal{P} \cap \mathcal{L}) \subseteq \mathcal{R}_k, \quad (3)$$

but equality may not hold in general.

We note that the entries of every matrix $A \in \mathcal{R}_k$ are in $[0, A_{\emptyset, \emptyset}]$. Indeed, looking at the 2×2 submatrix formed by the rows corresponding to some k -labeled flat graph F and the k -labeled edgeless graph U_k . From $A \in \Phi_k \mathcal{L}$ it follows that $A_{U_k, F} = A)F, F = A_{F, U_k}$, so positive semidefiniteness implies that $A_{U_k, U_k} A_{F, F} \geq A_{F, F}^2$. Since $A_{U_k, U_k} = A_{\emptyset, \emptyset}$ by $A \in \Phi_k \mathcal{L}$, we get that $(A_{\emptyset, \emptyset} - A_{F, F})A_{F, F} \geq 0$, which implies that $A_{F, F} \in [0, A_{\emptyset, \emptyset}]$.

For $k \leq m$, we consider \mathcal{F}_k as a subset of \mathcal{F}_m , by adding $m - k$ isolated nodes labeled $k + 1, \dots, m$. The corresponding restriction operator on matrices we denote by $\Phi_{m, k}$.

We claim that the following weak converse of (3) holds:

$$\Phi_k(\mathcal{P} \cap \mathcal{L}) = \bigcap_{m \geq k} \Phi_{m, k} \mathcal{R}_m. \quad (4)$$

Indeed, let A be a matrix that is contained in the right hand side. Then for every $m \geq k$ we have a matrix $B_m \in \mathcal{R}_m$ such that A is a restriction of B_m . Now let $m \rightarrow \infty$; by selecting a subsequence, we may assume that all entries of B_m tend to a limit. This limit defines a graph parameter f , which is normalized, isolate-indifferent and flatly reflection positive. By Proposition 3.4, f is reflection positive, and so the matrix $M(f)$ is in $\mathcal{P} \cap \mathcal{L}$ and $\Phi_k M(f) = A$.

Let $x = \alpha_1 F_1 + \dots + \alpha_r F_r$. We may assume that $|V(F_i)| = k$ for all i . Let F'_i be obtained from F_i by labeling all its nodes. Let $A \in \mathcal{M}_k$ denote the matrix

$$A_{FG} = \begin{cases} \alpha_i, & \text{if } F = G = F_i, \\ 0, & \text{otherwise.} \end{cases}$$

Then $x \geq 0$ means that $A \cdot Z \geq 0$ for all $Z \in \Phi_k(\mathcal{P} \cap \mathcal{L})$, in other words, A is in the dual cone of $\Phi_k(\mathcal{P} \cap \mathcal{L})$. From (4) it follows that there are diagonal matrices $A_m \in \mathcal{M}_k$ such that $A_m \rightarrow A$ and $A_m \cdot Y \geq 0$ for all $Y \in \Phi_{m, k} \mathcal{R}_m$. In other words, $A_m \cdot \Phi_{m, k} Z \geq 0$ for all $Z \in \mathcal{R}_m$, which can also be written as $\Phi_{m, k}^* A_m \cdot Z \geq 0$, where $\Phi_{m, k}^* : \mathcal{M}_k \rightarrow \mathcal{M}_m$ is the adjoint of the linear map $\Phi_{m, k} : \mathcal{M}_m \rightarrow \mathcal{M}_k$. (This adjoint acts by adding 0-s in all entries outside $\mathcal{F}_k \times \mathcal{F}_k$.) So $\Phi_{m, k}^* A_m$ is in the polar cone of $\mathcal{R}_m = \mathcal{P}_m \cap \mathcal{L}_m$, which is $\mathcal{P}_m^* + \mathcal{L}_m^*$. The positive semidefinite cone is self-polar. The linear space \mathcal{L}_m^* consists of those matrices $B \in \mathcal{M}_m$ for which $\sum_{F_1, F_2} B_{F_1, F_2} = 0$, where the summation extends over all pairs $F_1, F_2 \in \mathcal{F}'_m$ for which $F_1 F_2 \simeq F_0$ for some fixed graph F_0 . Thus we have $\Phi_{m, k}^* A_m = P + L$, where P is positive semidefinite and $L \in \mathcal{L}_m^*$. Since

P is positive semidefinite, we can write it as $P = \sum_{k=1}^N v_k v_k^\top$, where $v_k \in \mathbb{R}^{\mathcal{F}'_m}$. We can write this as

$$\sum_{\substack{F_1, F_2 \\ F_1 F_2 \simeq F_0}} \sum_{k=0}^N v_{k, F_1} v_{k, F_2} = \begin{cases} (A_m)_{F_0, F_0}, & \text{if } F_1 F_2 \simeq F_0 \in \mathcal{F}_k, \\ 0, & \text{otherwise.} \end{cases}$$

In other words,

$$\sum_{k=1}^N \left(\sum_F v_{k, F} F \right)^2 = \sum_{F_0} (A_m)_{F_0, F_0} F_0,$$

which proves the Theorem. \square

References

- [1] C. Borgs, J. Chayes, L. Lovász: Moments of Two-Variable Functions and the Uniqueness of Graph Limits, <http://arxiv.org/abs/0803.1244>
- [2] C. Borgs, J. Chayes, L. Lovász, V.T. Sós, K. Vesztergombi: Counting graph homomorphisms, in: *Topics in Discrete Mathematics* (ed. M. Klazar, J. Kratochvíl, M. Loeb, J. Matoušek, R. Thomas, P. Valtr), Springer (2006), 315–371.
- [3] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós, B. Szegedy and K. Vesztergombi: Graph Limits and Parameter Testing, Proc. 38th Annual ACM Symp. on Theory of Computing 2006, 261–270.
- [4] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós, and K. Vesztergombi: Convergent Graph Sequences I: Subgraph frequencies, metric properties, and testing, *Advances in Math.* (2008), 10.1016/j.aim.2008.07.008.
- [5] P. Erdős, L. Lovász, J. Spencer: Strong independence of graphcopy functions, in: *Graph Theory and Related Topics*, Academic Press, 165–172.
- [6] M. Freedman, L. Lovász, A. Schrijver: Reflection positivity, rank connectivity, and homomorphisms of graphs, *J. Amer. Math. Soc.* **20** (2007), 37–51.
- [7] J.B. Lasserre: A sum of squares approximation of nonnegative polynomials, *SIAM Review* **49** (2007), 651–669.
- [8] L. Lovász and B. Szegedy: Szemerédi’s Lemma for the analyst, *Geom. Func. Anal.* **17** (2007), 252–270.
- [9] L. Lovász, B. Szegedy: Limits of dense graph sequences, *J. Comb. Theory B* **96** (2006), 933–957.
- [10] A.F. Sidorenko: Inequalities for functionals generated by bipartite graphs (Russian) *Diskret. Mat.* **3** (1991), 50–65; translation in *Discrete Math. Appl.* **2** (1992), 489–504.
- [11] A.F. Sidorenko: A correlation inequality for bipartite graphs, *Graphs and Combin.* **9** (1993), 201–204.